

The Throughput of Slotted Aloha with Diversity

Andrea Munari, Michael Heindlmaier, Gianluigi Liva and Matteo Berioli

Abstract—In this paper, a simple variation of classical Slotted Aloha is introduced and analyzed. The enhancement relies on adding multiple receivers that gather different observations of the packets transmitted by a user population in one slot. For each observation, the packets transmitted in one slot are assumed to be subject to independent on-off fading, so that each of them is either completely faded, and then does not bring any power or interference at the receiver, or it arrives unfaded, and then may or may not, collide with other unfaded transmissions. With this model, a novel type of diversity is introduced to the conventional SA scheme, leading to relevant throughput gains already for moderate number of receivers. The analytical framework that we introduce allows to derive closed-form expression of both throughput and packet loss rate an arbitrary number of receivers, providing interesting hints on the key trade-offs that characterize the system. We then focus on the problem of having receivers forward the full set of collected packets to a final gateway using the minimum possible amount of resources, i.e., avoiding delivery of duplicate packets, without allowing any exchange of information among them. We derive what is the minimum amount of resources needed and propose a scheme based on random linear network coding that achieves asymptotically this bound without the need for the receivers to coordinate among them.

I. INTRODUCTION

A renewed interest for Aloha-like random access (RA) protocols led recently to the development of new high-throughput uncoordinated multiple-access schemes [1]–[10]. These schemes share the feature of cancelling the interference caused by a packet whenever (a portion of) it is successfully decoded. Among the aforementioned works, a specific class is based on the diversity slotted Aloha (DSA) protocol introduced in [11] enhanced by successive interference cancellation (SIC). In [4], [5] it was shown that the SIC process can be well modeled by means of a bipartite graph. By exploiting the graph model, a remarkably-high capacity (e.g., up to 0.8 [packets/slot]) can be achieved in practical implementations, whereas for large medium access control (MAC) frames it was demonstrated that fully efficiency (1 [packets/slot]) can be substantially attained [8], [10], [12]. A further key ingredient to attain large throughput gains deals with the exploitation of *diversity*. As an example, the approaches proposed in [1], [4]–[6], [9]–[11] take advantage of time diversity to resolve collisions.

In this paper, we develop and analyze a simple yet powerful *relay-aided* slotted Aloha (SA) scheme which enjoys space diversity. More specifically, K independent observations of

a slot are supposed to be available. The different observations are associated to K relays, and, for each of them, the transmitted packets are subject to independent fading coefficients. Collisions are regarded as destructive, and the system is complemented by having relays deliver what they have decoded to a centralized gateway.

SA with space (antenna) diversity was analyzed in [13] under the assumption of Rayleigh fading and shadowing, with emphasis on the two-antenna case. With respect to [13], we introduce in our analysis a simplified channel model. In particular, the uplink wireless link connecting user i and relay j is described by a packet erasure channel with erasure probability $\varepsilon_{i,j}$, following the on-off fading model [14]. The fading is assumed to be independent for each pair of user-relay pair. Despite its simplicity, the model is accurate enough for some cases of interest. As an example, it captures the main features of an interactive satellite network with satellite located on different orbits, and where the line-of-sight link between users and relays may be blocked whenever an obstacle lies between a user and a satellite (here, the satellites play the role of relays).

Under this fading model, elegant exact expressions for the system throughput as a function of the number of relays are derived, yielding deep insights in the gains provided by diversity in SA protocols. We further provide an analysis on how the link between the relays and the centralized gateways (also referred to as *downlink*) shall be dimensioned, assuming the relays to be uncoordinated. A bound on the downlink capacity is derived, which is achieved by a random linear coding approach based on Slepian-Wolf coding.

The rest of the paper is organized as follows. We start in Section II by defining the system model that is used to develop our framework. Section III provides a thorough analysis of the system uplink, characterizing it in terms of throughput and delivery reliability, whereas in Section IV we study how to effectively deliver collected packets to a common gateway without resorting to coordination and information exchange among relays. In Appendix, we also investigate, for the two-receiver case, an extension of the considered scheme that takes advantage of successive interference cancellation techniques.

II. SYSTEM MODEL AND PRELIMINARIES

Throughout this paper, we focus on the topology depicted in Fig. 1, where an infinite population of users want to deliver information in the form of data packets to a collecting gateway (GW). The transmission process is divided in two phases, referred to as *uplink* and *downlink*, respectively. During the former, data are sent in an uncoordinated fashion over a shared wireless channel to a set of K receivers or relays, which, in turn, forward collected information to the GW in the downlink.

Andrea Munari, Gianluigi Liva and Matteo Berioli are with Institute of Communication and Navigation of the Deutsches Zentrum für Luft- und Raumfahrt (DLR), D-82234 Weßling, Germany (e-mail: {andrea.munari,gianluigi.liva,matteo.beriolli}@dlr.de).

Michael Heindlmaier is with the Lehrstuhl für Nachrichtentechnik, Technische Universität München (TUM), D-80290 München Germany (e-mail: michael.heindlmaier@tum.de).




Fig. 1. Reference topology for the system under consideration.

As to the uplink, time is divided in successive slots, and transmission parameters in terms of packet length, coding and modulation are fixed such that one packet can be sent within one time unit. Users are assumed to be slot-synchronized, and Slotted Aloha (SA) [15] is employed as medium access policy. Furthermore, the number of users accessing the channel in a generic slot is modelled as a Poisson-distributed r.v. U of intensity ρ , with:

$$\Pr\{U = u\} = \frac{\rho^u e^{-\rho}}{u!}. \quad (1)$$

The uplink wireless link connecting user i and receiver j is described by a packet erasure channel with erasure probability $\varepsilon_{i,j}$, where independent realizations for any (i, j) pair as well as for a specific user-receiver couple across time slots are assumed. For the sake of mathematical tractability, we set $\varepsilon_{i,j} = \varepsilon$, $\forall i, j$. Following the on-off fading description [14], we assume that a packet is either completely shadowed, not bringing any power or interference contribution at a receiver, or it arrives unfaded. While, on the one hand, such a model is especially useful to develop mathematically tractable approaches to the aim of highlighting the key tradeoffs of the considered scenario, it also effectively captures effects like fading and short-term receiver unavailability due, for instance, to the presence of obstacles. Throughout our investigation, no multi-user detection capabilities are considered at the relays, so that collisions among non-erased data units are regarded as destructive and prevent decoding at a receiver.

Within this framework, the number of non-erased packets that arrive at a relay when u concurrent transmissions take place follows a binomial distribution of parameters $(u, 1 - \varepsilon)$ over one slot. Therefore, a successful reception occurs with probability $u(1 - \varepsilon)\varepsilon^{u-1}$, and the average throughput experienced at each of the K receivers, in terms of decoded packets per slot, can be computed as:

$$\mathcal{T}_{sa} = \sum_{u=0}^{\infty} \frac{\rho^u e^{-\rho}}{u!} u(1 - \varepsilon)\varepsilon^{u-1} = \rho(1 - \varepsilon)e^{-\rho(1 - \varepsilon)}, \quad (2)$$

corresponding to the performance of a SA system with erasures. On the other hand, a spatial diversity gain can

be triggered when the relays are considered jointly, since independent channel realizations may lead them to retrieve different information units over the same time slot. In order to quantify this beneficial effect, we label a packet as *collected* when it has been received by at least one of the relays, and we introduce the *uplink throughput* $\mathcal{T}_{up,K}$ as the average number of collected packets per slot. Despite its simplicity, such a definition offers an effective characterization of the beneficial effects of diversity, by properly accounting for both the possibility of retrieving up to $\min\{u, K\}$ distinct data units or multiple times the same data unit over a slot, as will be discussed in details in Section III-A. On the other hand, $\mathcal{T}_{up,K}$ also quantifies the actual amount of information that can be retrieved by the set of receivers, providing an upper bound for the overall achievable end-to-end performance, and setting the target for the design of any relay-to-GW delivery strategy.

For the downlink phase, we focus on a *decode and forward* (D&F) approach, so that each receiver re-encodes and transmits only packets it has correctly retrieved during the uplink phase, or possibly linear combinations thereof. A finite downlink capacity is assumed, and relays have to share a common bandwidth to communicate to the GW by means of a TDMA scheme. In order to get an insightful characterization of the optimum achievable system performance, we assume relay-to-GW links to be error free, and let resource allocation for the D&F phase be performed ideally and without additional cost by the central collecting unit.

We then complement our study in Appendix I by considering, for the simplified $K = 2$ scenario, an *amplify and forward* (A&F) approach. In this case, relays simply deliver an amplified version of the analog waveform (possibly the outcome of a collision) they received, whereas the GW performs decoding relying on successive interference cancellation (SIC) techniques. The goal of such an investigation is to derive a characterization of the gains that are achievable by jointly processing signals incoming at different receivers. Along this line of reasoning, we will focus on an idealized downlink, such that information can be reliably delivered to the collecting unit at no cost in terms of bandwidth.

A. Notation

Prior to delving into the details of our mathematical framework, we introduce in the following some useful notation. All the variables will be properly introduced when needed in the discussion, and the present section is simply meant to offer a quick reference point throughout the reading.

K relays are available, and, within time slot t , the countably infinite set of possible outcomes at each of them is labeled as $\Omega_t := \{\omega_0^t, \omega_1^t, \omega_2^t, \dots, \omega_\infty^t\}$ for each $t = 1, 2, \dots, n$. Here, ω_0^t denotes the erasure event (given either by a collision or by an idle slot), while ω_j^t indicates the event that the packet of the j -th user arriving in slot t was received. According to this notation, we define as X_k^t the random variables with alphabet \mathbb{N} , where $X_k^t = j$ if ω_j^t was the observation at relay k . When needed for mathematical discussion, we let the uplink operate for n time slots. In this case, let \mathcal{A}_k^n be the set of collected packets after n time slots at receiver k , where $\mathcal{A}_k^n \subsetneq$

$\bigcup_{t=1}^n \{\Omega_t \setminus \omega_0^t\}$. That is, we do not add the erasure events to \mathcal{A}_k^n . The number of received packets at relay k after n time slots is thus $|\mathcal{A}_k^n|$.

In general, the complement of a set \mathcal{A} is indicated as $\overline{\mathcal{A}}$. We write vectors as lowercase underlined variables, e.g., \underline{w} , while matrices and their transposes are labeled by uppercase letters, e.g., G and G^T .

III. A CHARACTERIZATION OF SYSTEM UPLINK

With reference to the topology of Fig. 1, we first consider the uplink phase. In order to gather a comprehensive description of the improvements enabled by receiver diversity, we characterize the system by means of two somewhat complementary metrics: uplink throughput (Section III-A) and packet loss rate (Section III-B).

A. Uplink Throughput

Let us focus on the random access channel, and, following the definition introduced in Section II, let C be the number of packets collected by the relays over one slot. C is a r.v. with outcomes in the set $\{0, 1, 2, \dots, K\}$, where the maximum value occurs when the K receivers decode distinct packets due to different erasure patterns. The average uplink throughput can thus be expressed by conditioning on the number of concurrent transmissions as:

$$\mathcal{T}_{up,K} = \mathbb{E}_U[\mathbb{E}[C|U]] = \sum_{u=0}^{\infty} \frac{\rho^u e^{-\rho}}{u!} \sum_{c=0}^K c \Pr\{C=c|U=u\}. \quad (3)$$

While Eq. (3) formula holds for any K , the computation of the collection probabilities intrinsically depends on the number of available relays. In this perspective, we articulate our analysis by first considering the two-receiver case, to then extend the results for an arbitrary topology.

1) *The Two-Receiver Case:* Let us first then focus on the case in which only two relays are available. Such a scenario allows a compact mathematical derivation of the uplink throughput, as the events leading to packet collection at the relays set can easily be expressed. On the other hand, it also represents a case of practical relevance, as it can be instantiated by simply adding a receiver to an existing SA-based system. When $K=2$, the situation for $C=1$ can easily be accounted for, since a single packet can be collected as soon as at least one of the relays does not undergo an erasure, i.e., with overall probability $1-\varepsilon^2$. On the other hand, by virtue of the binomial distribution of U , the event of collecting a single information unit over one slot occurs with probability

$$\Pr\{C=1|U=u\} = 2u(1-\varepsilon)\varepsilon^{u-1} [1-u(1-\varepsilon)\varepsilon^{u-1}] + u(1-\varepsilon)^2 \varepsilon^{2(u-1)}, \quad (4)$$

where the former addend accounts for the case in which one relay decodes a packet while the other does not (either due to erasures or to a collision), whereas the latter tracks the case of having the two relays decoding the same information unit. Conversely, a reward of two packets is obtained only when the receivers successfully retrieve distinct units, with probability

$$\Pr\{C=2|U=u\} = u(u-1)(1-\varepsilon)^2 \varepsilon^{2(u-1)}. \quad (5)$$




Fig. 2. Average uplink throughput vs channel load under different erasure probabilities. Black markers indicate the performance in the presence of two receivers, whereas white markers report the behavior of pure SA.

Plugging these results into (3) we get, after some calculations, a closed-form expression for the throughput in the uplink and thus, as discussed, also for the end-to-end D&F case with infinite downlink capacity:

$$\mathcal{T}_{up,2} = 2\rho(1-\varepsilon)e^{-\rho(1-\varepsilon)} - \rho(1-\varepsilon)^2 e^{-\rho(1-\varepsilon^2)}. \quad (6)$$

The trend of $\mathcal{T}_{up,2}$ is reported in Fig. 2 against the channel load ρ for different values of the erasure probability, and compared to the performance in the presence of a single receiver, i.e., \mathcal{T}_{sa} . Eq. (6) conveniently expresses $\mathcal{T}_{up,2}$ as twice the throughput of SA in the presence of erasures, reduced by a loss factor which accounts for the possibility of having both relays decode the same information unit. In this perspective, it is interesting to evaluate the maximum throughput $\mathcal{T}_{up,2}^*(\varepsilon)$ as well as the optimal working point $\rho^*(\varepsilon)$ achieving it for the system uplink. The transcendental nature of (6) does not allow to obtain a closed formulation of these quantities, which, on the other hand, can easily be estimated by means of numerical optimization techniques. The results of this analysis are reported in Fig. 3, where the peak throughput \mathcal{T}_{up}^* is depicted by the black curve as a function of ε and compared to the performance of SA, which clearly collects on average at most 0.36 pkt/slot regardless of the erasure rate. In ideal channel conditions, i.e., $\varepsilon=0$, no benefits can be obtained by resorting to multiple relays, as all of them would see the same reception set across slots. Conversely, higher values of ε favour a decorrelation of the pattern of packets that can be correctly retrieved, and consequently improve the achievable throughput at the expense of higher loss rates. The result is a monotonically increasing behavior for $\mathcal{T}_{up,2}^*(\varepsilon)$, prior to plummeting with a singularity to a null throughput for the degenerate case $\varepsilon=1$. Fig. 3 also reports (circled-white markers) the average throughput obtained for $\rho=1/(1-\varepsilon)$, i.e., when the uplink of the system under consideration operates at the optimal working point for a single-receiver SA, showing a tight match. In fact, even though the abscissa of maximum $\rho^*(\varepsilon)$ may differ from this value (they coincide only for the ideal case $\varepsilon=0$), the error which




Fig. 3. Maximum uplink throughput vs erasure rate. The black continuous line reports the performance $\mathcal{T}_{up,2}^*$ of a two-receiver scheme, while white circled markers indicate $\mathcal{T}_{up,2}(1/(1-\varepsilon))$, and the dotted line shows the behavior of pure SA. Gray curves and markers are to be referred to the *amplify and forward* case, that will be treated in Appendix I.

is committed when approximating $\mathcal{T}_{up,2}^*$ with $\mathcal{T}_{up,2}(1/(1-\varepsilon))$ can easily be shown numerically to never exceed 0.6%, due to the very small slope of the function in the neighborhood of $\rho^*(\varepsilon)$. We can thus provide a very precise estimate of the peak uplink performance for a specific erasure rate as:

$$\mathcal{T}_{up,2}^*(\varepsilon) \simeq \frac{2}{\varepsilon} - (1-\varepsilon)e^{-1-\varepsilon}, \quad 0 \leq \varepsilon < 1. \quad (7)$$

which once again compactly captures the behavior of a two-receiver scenario by quantifying the loss with respect to twice the performance of SA. In this perspective, two remarks shall be made. First of all, in order to approach the upper bound, the system has to be operated at very high load, as $\rho^* \simeq 1/(1-\varepsilon)$. These working points are typically not of interest, since very low levels of reliability can be provided by a congested channel with high erasure rates. Nevertheless, the presence of a second receiver triggers remarkable improvements already for loss probabilities that are of practical relevance, e.g., under harsh fading conditions or for satellite networks. Indeed, with $\varepsilon = 0.1$ a $\sim 15\%$ raise can be spotted, whereas a loss rate of 20% already leads to a 50% throughput gain. Secondly, the proposed framework highlights how no modifications in terms of system load are needed with respect to plain SA for a two-receiver system to be very efficiently operated. Such a result is particularly interesting, as it suggests that a relay node can be seamlessly and efficiently added to an already operating SA uplink when available, triggering the maximum achievable benefit without the need to undergo a re-tuning of the system which might be particularly expensive in terms of resources.

2) *The General Case, $K > 2$* : Let us now focus on the general topology reported in Fig. 1, where K relays are available. While conceptually applicable, the approach presented to compute the uplink throughput in the two-receiver case becomes cumbersome as K grows, due to the rapidly increasing number of events that have to be accounted for. In order to characterize $\mathcal{T}_{up,K}$, then, we follow a different strategy. With

reference to a single slot t , let $\Omega_t := \{\omega_0^t, \omega_1^t, \omega_2^t, \dots, \omega_\infty^t\}$ for each $t = 1, 2, \dots, n$ be the countably infinite set of possible outcomes at each relay, where ω_0^t denotes the erasure event while ω_j^t indicates the event that the packet of the j -th user arriving in slot t was received. Let us furthermore define as X_k^t the random variables with alphabet $\mathcal{X} = \{0, 1, 2, \dots, \infty\}$, where $X_k^t = j$ if ω_j^t was the observation at relay k , so that $X_k^1, X_k^2, \dots, X_k^n$ is an i.i.d. sequence for each relay k . We let the uplink operate for n time slots, and indicate as \mathcal{A}_k^n the set of packets collected at receiver k over this time-span, where $\mathcal{A}_k^n \subseteq \bigcup_{t=1}^n \{\Omega_t \setminus \omega_0^t\}$ (i.e., we do not add the erasure events to \mathcal{A}_k^n). The number of received packets at relay k after n time slots is thus $|\mathcal{A}_k^n|$ and, with reference to this notation, we prove the following result:

Proposition 1: For an arbitrary number of K relays, the throughput $\mathcal{T}_{up,K}$ is given by

$$\mathcal{T}_{up,K} = \sum_{k=1}^K (-1)^{k-1} \binom{K}{k} \rho (1-\varepsilon)^k e^{-\rho(1-\varepsilon^k)} \quad (8)$$

Proof: We have $|\mathcal{A}_k^n| = \sum_{t=1}^n \mathbb{1}_{\{X_k^t \neq 0\}}$, where $\mathbb{1}_{\{E\}}$ denotes the indicator random variable that takes on the value 1 if the event E is true and 0 otherwise. The throughput seen by a single relay can then be written as $\mathcal{T}_{up,1} = \mathbb{E}[\mathbb{1}_{\{X_k^t \neq 0\}}] = \Pr\{X_k^t \neq 0\}$, and does not depend on the specific receiver being considered. By the weak law of large numbers,

$$\mathcal{T}_{up,1} = \lim_{n \rightarrow \infty} \frac{|\mathcal{A}_k^n|}{n} \quad (9)$$

or, more formally,

$$\lim_{n \rightarrow \infty} \Pr \left\{ \left| \frac{|\mathcal{A}_k^n|}{n} - \mathcal{T}_{up,1} \right| > \epsilon \right\} = 0 \text{ for some } \epsilon > 0. \quad (10)$$

Similarly, for K relays we have

$$\mathcal{T}_{up,K} = \lim_{n \rightarrow \infty} \frac{|\bigcup_{k=1}^K \mathcal{A}_k^n|}{n} \quad (11)$$

By the inclusion-exclusion principle (see, e.g., [16]), we have

$$\left| \bigcup_{k=1}^K \mathcal{A}_k^n \right| = \sum_{S \subseteq \{1, \dots, K\}, S \neq \emptyset} (-1)^{|S|-1} |\mathcal{I}_S^n| \quad (12)$$

$$\text{with } \mathcal{I}_S^n = \bigcap_{k \in S} \mathcal{A}_k^n \quad (13)$$

Here, \mathcal{I}_S^n denotes the set of packets that all the relay nodes specified by $S = \{k_1, k_2, \dots, k_{|S|}\}$ have in common:

$$|\mathcal{I}_S^n| = \left| \bigcap_{k \in S} \mathcal{A}_k^n \right| = \sum_{t=1}^n \mathbb{1}_{\{0 \neq X_{k_1}^t = X_{k_2}^t = \dots = X_{k_{|S|}}^t\}} \quad (14)$$

Due to symmetry in the setup, the value of $|\mathcal{I}_S^n|$ only depends on the cardinality of S but not the explicit choice, so that $|\mathcal{I}_S^n| = a_k^n$ for $k = |S|$, and,

$$\left| \bigcup_{k=1}^K \mathcal{A}_k^n \right| = \sum_{k=1}^K (-1)^{k-1} \binom{K}{k} a_k^n.$$

As $X_k^1, X_k^2, \dots, X_k^n$ are i.i.d., by the weak law of large numbers we have:

$$\lim_{n \rightarrow \infty} \frac{|\mathcal{I}_S^n|}{n} = \Pr[\{0 \neq X_{k_1}^t = X_{k_2}^t = \dots = X_{k_{|S|}}^t\}]. \quad (15)$$




Fig. 4. Average uplink throughput vs channel load for different number of relays K . The erasure probability has been set to $\varepsilon = 0.5$.

We can compute the latter probability as

$$\begin{aligned}
 & \Pr\{0 \neq X_{k_1}^t = \dots = X_{k_{|S|}}^t\} \\
 &= \sum_u \Pr\{0 \neq X_{k_1}^t = \dots = X_{k_{|S|}}^t | U = u\} \Pr\{U = u\} \\
 &= \sum_{u=1}^{\infty} \frac{e^{-\rho} \rho^u}{u!} \binom{u}{1} ((1-\varepsilon)\varepsilon^{u-1})^{|S|} \\
 &= (1-\varepsilon)^{|S|} \rho e^{-\rho(1-\varepsilon)^{|S|}}
 \end{aligned} \tag{16}$$

As $\lim_{n \rightarrow \infty} \frac{a_n^n}{n} = (1-\varepsilon)^k \rho e^{-\rho(1-\varepsilon)^k}$, the proposition follows. ■

The performance achievable by increasing the number of relays is reported against the channel load in Fig. 4 for a reference erasure rate $\varepsilon = 0.2$. As expected, $\mathcal{T}_{up,K}$ benefits from a higher degree of spatial diversity, showing how the system can collect more than one packet per uplink slot as soon as more than four receivers are available, for the parameters under consideration. Such a result stems from two main factors. On the one hand, increasing K enables larger peak throughput over a single slot, as up to K different data units can be simultaneously retrieved. On the other hand, broader receivers sets improve the probability of decoding packets in the presence of collisions even when less than K users accessed the channel, by virtue of the different erasure patterns they experience. The uplink throughput characterization is complemented by Fig. 5, which reports the peak value for $\mathcal{T}_{up,K}^*$ (solid black curve), obtained by properly setting the channel load to ρ_K^* (whose values are shown by the gray dashed curve), for an increasing relay population.¹ The plot clearly highlights how the benefit brought by introducing an additional receiver to the scheme, quantified by Eq. (17), progressively reduces, leading to a growth rate for the achievable throughput that is less than

¹ As discussed for the $K = 2$ case, a mathematical derivation of the optimal working point load ρ_K^* is not straightforward, and simple numerical maximization techniques were employed to obtain the results of Fig 5.




Fig. 5. Maximum achievable throughput $\mathcal{T}_{up,K}^*$ as a function of the number of relays K for an erasure rate $\varepsilon = 0.5$. The gray curve reports the load on the channel ρ_K^* needed to reach $\mathcal{T}_{up,K}^*$.

linear and that exhibits a logarithmic-like trend in K .

$$\begin{aligned}
 \Delta \mathcal{T}_{up} &= \mathcal{T}_{up,K} - \mathcal{T}_{up,K-1} \\
 &= \sum_{k=1}^K (-1)^{k-1} \binom{K-1}{k-1} \rho (1-\varepsilon)^k e^{\rho(1-\varepsilon)^k}
 \end{aligned} \tag{17}$$

B. Packet loss probability

The aggregate throughput derived in Section III-A represents a metric of interest towards understanding the potential of SA with diversity when aiming at reaping the most out of uplink bandwidth. On the other hand, operating an Aloha-based system at the optimal load ρ_K^* exposes each transmitted packet to a loss probability that may not be negligible. In the classical single-receiver case without fading, for instance, the probability for a data unit not to be collected evaluates to $1 - e^{-1} \simeq 0.63$. From this standpoint, in fact, several applications may resort to a lightly loaded random access uplink, aiming at a higher level of delivery reliability rather than at a high throughput. This is the case, for example, of channels used for logon and control signalling in many practical wireless networks. In order to investigate how diversity can improve performance in this direction, we extend our framework by computing the probability ζ_K that a user accessing the channel experiences a data loss, i.e., that the information unit it sends is not collected, either due to fading or to collisions, by any of the K relays.

To this aim, let \mathcal{O} describe the event that the packet of the observed user sent over time slot t is not received by any of the receivers. Conditioning on the number of interferers i , i.e., of data units that were concurrently present on the uplink channel at t , the sought probability can be written as:

$$\zeta_K = \sum_{i=0}^{\infty} \Pr[\mathcal{O} | I = i] \Pr[I = i]. \tag{18}$$

Here, the conditional probability can easily be determined recalling that each of the K relays experiences an independent




Fig. 6. Probability ζ_K that a packet sent by a user is not received by any of the relays. Different curves indicate different values of K , while the erasure probability has been set to $\varepsilon = 0.2$.

erasure pattern, obtaining $\Pr[\mathcal{O}|I = i] = (1 - (1 - \varepsilon)\varepsilon^i)^K$ for an individual packet and K relays with independent erasures on all individual links. By resorting to the binomial theorem, such an expression can be conveniently reformulated as:

$$\Pr[\mathcal{O}|I = i] = \sum_{k=0}^K (-1)^k \binom{K}{k} ((1 - \varepsilon)\varepsilon^i)^k. \quad (19)$$

On the other hand, the number of interferers seen by a user that accesses the channel at time t still follows a Poisson distribution of intensity ρ , so that, after simple calculations we finally get:

$$\zeta_K = \sum_{k=0}^K (-1)^k \binom{K}{k} (1 - \varepsilon)^k e^{-\rho(1 - \varepsilon^k)}. \quad (20)$$

Fig. 6 reports the behavior of ζ_K as a function of ρ when the erasure rate over a single link is set to $\varepsilon = 0.2$. Different lines indicate the trend when increasing the number of receivers from 1 to 10. As expected, when $\rho \rightarrow 0$, a user accessing the channel is not likely to experience any interference, so that failures can only be induced by erasures, leading to an overall loss probability of ε^K . In this perspective, the availability of multiple receivers triggers a dramatic improvement, enabling levels of reliability that would otherwise not be possible irrespective of the channel configuration. On the other hand, Eq. (20) turns out to be useful for system design, as it allows to determine the load that can be supported on the uplink channel while guaranteeing a target loss rate. Also in this case diversity can significantly ameliorate the performance. As shown in Fig. 6, for example, a target loss rate $\zeta = 5 \cdot 10^{-2}$ is achieved by a three- and four-receiver scheme under 6- and 10-fold larger loads compared to the $K = 2$ case, respectively.

IV. DOWNLINK STRATEGIES

The analysis carried out in Section III has characterized the average number of packets that can be decoded at the relay set when SA is used in the uplink. We now instead consider the complementary task of delivering what has been

collected to a central GW. In doing so, we aim at employing the minimum number of resources in terms of transmissions that have to be performed by the relays, while not allowing any information exchange among them. In particular, we consider a finite-capacity downlink, where the K receivers share a common bandwidth to communicate with the GW by means of a time division multiple access scheme, and we assume that each of them can reliably deliver exactly one packet, possibly composed of a linear combination of what has been collected, over one time unit. We once again focus on a horizon of n slots to operate the uplink, after which the downlink phase starts.

We structure our analysis in two parts. First, in Section IV-A, we derive lower bounds for the rates (in terms of downlink slots allocated per uplink slot) that have to be assigned to receivers in order to deliver the whole set of data units collected in the uplink over the n slots. Then, Section IV-B shows how a simple forwarding strategy based on random linear network coding suffices to achieve optimality, completing the downlink phase in $\mathcal{T}_{up,K}$ slots for asymptotically large values of n .

Prior to delving into the details, let us introduce some useful notation. We denote the L -bit data part of packet of the j -th user arriving in time slot t as $W_k^t \in \mathcal{W}$, with $\mathcal{W} = \mathbb{F}_{2^L} \cup e$, where e is added as the erasure symbol. We furthermore assume that the receiver can determine the corresponding user through a packet header, i.e., the receiver knows both j and t after successful reception. As the uplink operates over n time slots, relay k observes the vector $\underline{w}_k = [W_k^1, \dots, W_k^n]$. In each time slot, the tuple $(W_1^t, W_2^t, \dots, W_K^t)$ is drawn from a joint probability distribution $P_{W_1 \dots W_K}$ which is governed by the uplink, and different relays might receive the same packet.

A. Bounds for Downlink Rates

Each relay k transmits a packet in each of its nR_k downlink slots. We are interested in the set of rates $(R_k)_{k=1}^K$ such that the gateway can recover all packets (with high probability).

This is essentially the problem of distributed source coding (SW-Coding [17]), with the following modification: SW-coding ensures that the gateway can recover all K observed strings \underline{w}_k , $k = 1, 2, \dots, K$ perfectly. In this setup, the gateway should be able to recover every packet that was received at any relay. However, neither is the gateway interested in erasures symbols at the relays, i.e. whenever $W_k^t = e$ for any k, t , nor in reconstructing each relay sequence perfectly. The authors in [18] overcame this problem by assuming that the decoder knows all the erasure positions of the whole network. This assumption applies in our case as packet numbers are supposed to be known via a packet header. Let all erasure positions be represented by Γ .

The rates (R_1, \dots, R_K) are achievable [17] if

$$\sum_{k \in S} R_k \geq H(W_S | W_{\bar{S}}, \Gamma), \quad \forall S \subseteq [1, 2, \dots, K] \quad (21)$$

where $W_S = (W_{k_1}^t, W_{k_2}^t, \dots, W_{k_{|S|}}^t)$, denotes the observations at some time t at the subset of receivers specified by $S = \{k_1, k_2, \dots, k_{|S|}\}$. Γ has the effect of removing the

influence of the erasure symbols on the conditional entropies. Computing the entropies however requires the full probability distribution $P_{W_1 \dots W_K}$ which is a difficult task in general. By different means, we can obtain the equivalent conditions:

Proposition 2: The rates $(R_k)_{k=1}^K$ have to satisfy

$$\sum_{k \in \mathcal{S}} R_k \geq \mathcal{T}_{up,K} + \sum_{k=1}^{K-|\mathcal{S}|} (-1)^k \binom{K-|\mathcal{S}|}{k} \rho(1-\varepsilon)^k e^{-\rho(1-\varepsilon^k)}, \quad \forall \mathcal{S} \subseteq \{1, \dots, K\} \quad (22)$$

Proof: Consider a subset of relays $\mathcal{S} \subseteq \{1, \dots, K\}$ and their buffer contents $\bigcup_{k \in \mathcal{S}} \mathcal{A}_k^n$ after n time slots. In order to satisfy successful recovery at the gateway, at least all packets that have been collected only by nodes in the set \mathcal{S} and not by anyone else have to be communicated to the gateway. That is,

$$\sum_{k \in \mathcal{S}} n \cdot R_k \geq \left| \bigcup_{k \in \mathcal{S}} \mathcal{A}_k^n \setminus \bigcup_{k \in \overline{\mathcal{S}}} \mathcal{A}_k^n \right|, \quad (23)$$

with $\overline{\mathcal{S}} = \{1, \dots, K\} \setminus \mathcal{S}$. Note that

$$\bigcup_{k \in \mathcal{S}} \mathcal{A}_k^n \setminus \bigcup_{k \in \overline{\mathcal{S}}} \mathcal{A}_k^n = \mathcal{A}^n \setminus \bigcup_{k \in \overline{\mathcal{S}}} \mathcal{A}_k^n,$$

so by the inclusion-exclusion principle and due to $|\overline{\mathcal{S}}| = K - |\mathcal{S}|$

$$\left| \bigcup_{k \in \mathcal{S}} \mathcal{A}_k^n \setminus \bigcup_{k \in \overline{\mathcal{S}}} \mathcal{A}_k^n \right| = |\mathcal{A}_K^n| + \sum_{k=1}^{K-|\mathcal{S}|} (-1)^k \binom{K-|\mathcal{S}|}{k} a_k^n \quad (24)$$

with $|\mathcal{I}_S^n| = a_k^n$ for $k = |\mathcal{S}|$ as before. By plugging in the value for $\lim_{n \rightarrow \infty} \frac{a_k^n}{n}$, the proposition follows. ■

B. Random Linear Coding

By means of Proposition 2, we have derived a characterization of the rates that have to be assigned to relays in order to deliver the whole set of collected packets to the GW. In this section, we complete the discussion by proposing a strategy that is capable of matching such conditions, thus achieving optimality. The solution that we employ is based on a straightforward application of the well-known random linear coding scheme in [19], and will therefore only briefly sketched in the following.

Each relay k generates a matrix $G_k \in \mathbb{F}_{2L}^{nR_k \times n}$ and obtains the data part of its nR_k transmit packets by $\underline{c}_k^T = G_k \underline{w}_k^T$. Whenever an element of \underline{w}_k was an erasure symbol, the corresponding column of G_k is an all-zero column. Erasure symbols thus have no contributions to the transmit packets \underline{c}_k . All other elements of G_k are drawn uniformly at random from \mathbb{F}_{2L}^* , where \mathbb{F}_{2L}^* denotes the multiplicative group of \mathbb{F}_{2L} .

The gateway collects all incoming packets and obtains the system of linear equations

$$\underbrace{\begin{pmatrix} \underline{c}_1^T \\ \underline{c}_2^T \\ \vdots \\ \underline{c}_K^T \end{pmatrix}}_{\underline{c}^T} = \underbrace{\begin{pmatrix} G_1 & 0 & \dots & 0 \\ 0 & G_2 & \dots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \dots & G_K \end{pmatrix}}_G \underbrace{\begin{pmatrix} \underline{w}_1^T \\ \underline{w}_2^T \\ \vdots \\ \underline{w}_K^T \end{pmatrix}}_{\underline{w}^T} \quad (25)$$

where $G \in \mathbb{F}_{2L}^{n \sum_k R_k \times nK}$. Note that some elements of \underline{w} can be identical because they were received by more than one relay and thus are elements of some $\underline{w}_{k_1}, \underline{w}_{k_2}, \dots$. One can merge these entries in \underline{w} that appear more than once. Additionally, we drop all erasure-symbols in \underline{w} and delete the corresponding columns in G to obtain the reduced system of equations

$$\underline{c}^T = \tilde{G} \tilde{w}^T \quad (26)$$

where $\tilde{w} \in \mathbb{F}_{2L}^{|\mathcal{A}^n|}$ contains only distinct received packets and no erasure symbols. Clearly, there are $|\bigcup_{k \in \mathcal{S}} \mathcal{A}_k^n|$ elements in \tilde{w} .

We partition the entries in \tilde{w} into $2^K - 1$ vectors \tilde{w}_S for each nonempty subset $S \subseteq \{1, 2, \dots, K\}$: Each vector \tilde{w}_S contains all packets that have been received only by all relays specified by S and not by anyone else. That is, \tilde{w}_S corresponds to the set $\mathcal{P}_S^n = \bigcap_{k \in \mathcal{S}} \mathcal{A}_k^n \setminus \bigcup_{k \in \overline{\mathcal{S}}} \mathcal{A}_k^n$, its length is $|\mathcal{P}_S^n|$.

The columns in \tilde{G} and rows in \tilde{w}^T can be permuted such that one can write

$$\underline{c}_k^T = \tilde{G}_k \tilde{w} := \sum_{S \subseteq \{1, 2, \dots, K\}} \tilde{G}_{k,S} \cdot \tilde{w}_S, \quad \forall k = 1, \dots, K. \quad (27)$$

Each of the matrices $\tilde{G}_{k,S} \in \mathbb{F}_{2L}^{nR_k \times |\mathcal{P}_S^n|}$ contains only elements from \mathbb{F}_{2L}^* if $k \in S$ and is an all-zero matrix otherwise. A compact representation for $K = 3$ is shown in (28) at the bottom of next page.

The variables that are involved only in $n \sum_{k \in \mathcal{S}} R_k$ equations are those in $\tilde{w}_{\mathcal{L}}$, $\mathcal{L} \subseteq \mathcal{S}$, for each subset $\mathcal{S} \subseteq \{1, 2, \dots, K\}$. For decoding, the number of equations has to be larger or equal to the number of variables, so a necessary condition for decoding is that $n \sum_{k \in \mathcal{S}} R_k \geq \sum_{\mathcal{L} \subseteq \mathcal{S}} |\mathcal{P}_{\mathcal{L}}^n|$. This is satisfied by (23), since $\sum_{\mathcal{L} \subseteq \mathcal{S}} |\mathcal{P}_{\mathcal{L}}^n| = |\bigcup_{k \in \mathcal{S}} \mathcal{A}_k^n \setminus \bigcup_{k \in \overline{\mathcal{S}}} \mathcal{A}_k^n|$, as we show in Appendix II

A sufficient condition is that the matrix \tilde{G}_k , $k \in \mathcal{S}$ representing $n \sum_{k \in \mathcal{S}} R_k$ equations has rank $\sum_{\mathcal{L} \subseteq \mathcal{S}} |\mathcal{P}_{\mathcal{L}}^n|$ for each subset $\mathcal{S} \subseteq \{1, 2, \dots, K\}$. Denote the set of indices of nonzero columns of matrix \tilde{G}_k as the support of \tilde{G}_k . Note that a row of matrix \tilde{G}_k has a different support than a row of matrix \tilde{G}_l , for $k \neq l$. These rows are thus linearly independent. It thus suffices to check that all rows of matrix \tilde{G}_k are linearly independent. As all nonzero elements are randomly drawn from \mathbb{F}_{2L}^* , the probability of linear dependence goes to zero as L grows large, completing our proof, and showing that the presented forwarding scheme achieves the bounds of Proposition 2.

V. CONCLUSIONS

In this paper, a simple and practical extension of Slotted Aloha in the presence of multiple receivers, or relays, has been presented and thoroughly discussed. By means of an analytical framework, closed-form expression for the uplink throughput (defined as the average number of packets per slot *collected* by the set of K relays), as well as for the probability that a data unit is not retrieved by any of the receivers, have been derived for an arbitrary value of K under the assumption of on-off fading. Remarkable gains have been shown and discussed already for a moderate number of receivers. The study is complemented by considering the problem of delivering the set of collected packets to a common gateway without allowing any information exchange among receivers. Theoretical bounds for the amount of resources that have to be allocated to achieve this task have been derived, and a simple scheme based on random linear network coding has been shown to match such bounds.

APPENDIX I

AMPLIFY AND FORWARD WITH SIC AT THE GATEWAY

The framework developed in this paper has focused on characterizing the performance achievable by a SA system with receiver diversity when a D&F scheme is implemented at intermediate nodes. On the other hand, restricting relays to simply forward what they have successfully retrieved in the uplink prevents the GW from performing joint decoding on possibly uncorrelated signals. In order to go beyond this limitation, we consider in this appendix the possibility for receivers to send in the downlink, on a slot-basis, an amplified version of the analog waveform they perceive even in the presence of a collision, following an *amplify and forward* (A&F) approach. For the sake of mathematical tractability, we focus on the $K = 2$ case, and we model relays to instantly and reliably deliver information to the GW. The advantage of such an assumption is twofold. On the one hand, it will allow us once more to identify elegant closed-form expressions for the throughput of A&F with spatial diversity, highlighting the fundamental tradeoffs that arise in the presence of multiple receivers. On the other hand, despite its ideality, the model under consideration is representative for several scenarios of practical interest, in which the bandwidth available in the downlink is much larger than the one of the uplink. Satellite networks, as well as topologies where multiple base stations or access points are connected to a coordinating unit via a wideband backbone may be examples in this direction.

At the gateway side, successive interference cancellation (SIC) is applied to the collected signals. This approach offers an improvement whenever the waveform forwarded by one of the relays allows decoding of a packet, say x , while the other reports a collision given by the superposition of x and one other packet. In this condition, the set of relays would be able to collect only one information unit, whereas, with A&F, the gateway can subtract the interference contribution of x from the collision-corrupted waveform and successfully collect the second packet as well. Details on the accuracy of this model on noisy channels with actual signal processing techniques can be found in [1], [4].

The gain offered by SIC can thus be computed for the $K = 2$ case by simply adding to the uplink throughput derived in Section III-A one additional collected data unit each time the described collision condition is met. Hence, we can write:

$$\mathcal{G}_{SIC} = \sum_{u=2}^{\infty} \frac{\rho^u e^{-\rho}}{u!} 2u(u-1) \varepsilon^{u-1} (1-\varepsilon) \cdot \varepsilon^{u-2} (1-\varepsilon)^2 = 2\rho^2 \varepsilon (1-\varepsilon)^3 e^{-\rho(1-\varepsilon^2)}, \quad (29)$$

where, within the summation, $\varepsilon^{u-1}(1-\varepsilon)$ accounts for the correct reception at one relay while $\varepsilon^{u-2}(1-\varepsilon)^2$ enforces a collision of exactly two packets at the other relay, for a total of $2u(u-1)$ configurations that can be solved with SIC. The average number of collected packets at the GW per uplink slot, which we refer to as $\mathcal{T}_{A\&F}$, is thus simply expressed as $\mathcal{T}_{A\&F} = \mathcal{T}_{up,2} + \mathcal{G}_{SIC}$:

$$\mathcal{T}_{A\&F} = 2\rho(1-\varepsilon) e^{-\rho(1-\varepsilon)} - \rho(1-\varepsilon)^2 e^{-\rho(1-\varepsilon^2)} + 2\rho^2 \varepsilon (1-\varepsilon)^3 e^{-\rho(1-\varepsilon^2)}. \quad (30)$$

The obtained trend is plotted in Fig. 7 against the channel load ρ for an erasure probability of $\varepsilon = 0.2$, showing a 20% and 66% improvement in peak throughput compared to the performance achieved under the same conditions in the uplink (i.e., without SIC) and by a SA scheme with single receiver, respectively. As discussed in Section III-A, a closed-form evaluation of the maximum throughput $\mathcal{T}_{A\&F}^*$ is not straightforward, due to the transcendental nature of the terms that define the metric. Nevertheless, in the two-receiver case, a good approximation is once again offered by evaluating $\mathcal{T}_{A\&F}$ at $\rho = 1/(1-\varepsilon)$, obtaining, after some calculations:

$$\mathcal{T}_{A\&F}^*(\varepsilon) \simeq \frac{2}{e} - e^{-1-\varepsilon} (1-3\varepsilon+2\varepsilon^2), \quad (31)$$

where a loss factor of $e^{-1-\varepsilon}(1-\varepsilon)(1-2\varepsilon) \leq e^{-1-\varepsilon}(1-\varepsilon)$ is exhibited with respect to the upper bound provided by twice

$$\begin{pmatrix} \underline{c}_1^T \\ \underline{c}_2^T \\ \underline{c}_3^T \end{pmatrix} = \begin{pmatrix} \tilde{G}_1 \\ \tilde{G}_2 \\ \tilde{G}_3 \end{pmatrix} \underline{\tilde{w}}^T = \begin{pmatrix} \tilde{G}_{1,\{1\}} & 0 & 0 & \tilde{G}_{1,\{1,2\}} & \tilde{G}_{1,\{1,3\}} & 0 & \tilde{G}_{1,\{1,2,3\}} \\ 0 & \tilde{G}_{2,\{2\}} & 0 & \tilde{G}_{2,\{1,2\}} & 0 & \tilde{G}_{2,\{2,3\}} & \tilde{G}_{2,\{1,2,3\}} \\ 0 & 0 & \tilde{G}_{3,\{3\}} & 0 & \tilde{G}_{3,\{1,3\}} & \tilde{G}_{3,\{2,3\}} & \tilde{G}_{3,\{1,2,3\}} \end{pmatrix} \begin{pmatrix} \underline{\tilde{w}}_{\{1\}}^T \\ \underline{\tilde{w}}_{\{2\}}^T \\ \underline{\tilde{w}}_{\{3\}}^T \\ \underline{\tilde{w}}_{\{1,2\}}^T \\ \underline{\tilde{w}}_{\{1,3\}}^T \\ \underline{\tilde{w}}_{\{2,3\}}^T \\ \underline{\tilde{w}}_{\{1,2,3\}}^T \end{pmatrix} \quad (28)$$




Fig. 7. End-to-end throughput vs uplink channel load with infinite downlink capacity and $\varepsilon = 0.2$. The black-marked curve reports the behavior with A&F; the gray-marked curve indicates D&F; the white-marked curve shows the behavior of SA with a single receiver.

the throughput of SA. The behavior of $\mathcal{T}_{A\&F}^*(\varepsilon)$ is reported in Fig. 3, where the dashed-gray curve indicates the actual peak throughput values computed numerically, whereas the white-squared markers report the proposed approximation. The plot highlights how, as opposed to what discussed for the non-SIC case, the introduction of joint processing modifies the shape of the curve, identifying an optimal (albeit not practical for many applications) erasure probability where the throughput is more than doubled over SA. On the other hand, it is remarkable to point out that the most relevant improvements over a non-SIC multi-receiver uplink are triggered exactly for values of ε that may indeed be experienced in practical scenarios, boosting the peak throughput by up to 25%.

As a concluding observation, notice how the two-relay A&F solution that we discussed requires rather simple interference cancellation procedures compared to other advanced random access schemes [1], [4], as only two observations need to be considered for joint decoding. The presented architecture, thus, represents an interesting tradeoff between complexity and performance gain, and triggers interest in more advanced scenarios where more receivers are available.

APPENDIX II

We derive that $\sum_{\mathcal{L} \subseteq \mathcal{S}} |\mathcal{P}_{\mathcal{L}}^n| = |\bigcup_{k \in \mathcal{S}} \mathcal{A}_k^n \setminus \bigcup_{k \in \mathcal{S}} \mathcal{A}_k^n|$. We need the following lemma.

Lemma 1: For a collection of sets $\mathcal{B}_1, \mathcal{B}_2, \dots, \mathcal{B}_K$ and a subset $\mathcal{S} \subseteq \{1, \dots, K\}$,

$$\bigcup_{k \in \mathcal{S}} \mathcal{B}_k = \bigcup_{\mathcal{L} \subseteq \mathcal{S}} \left(\bigcap_{l \in \mathcal{L}} \mathcal{B}_l \setminus \bigcup_{s \in \mathcal{S} \setminus \mathcal{L}} \mathcal{B}_s \right), \quad (32)$$

where the sets on the RHS do not intersect and thus form a partition of the LHS.

Proof: We first show that any element $b \in \bigcup_{k \in \mathcal{S}} \mathcal{B}_k$ is also included in the RHS: Pick an element $b \in \bigcup_{k \in \mathcal{S}} \mathcal{B}_k$. Assume \mathcal{L} is the subset of largest cardinality such that $b \in \mathcal{B}_l$, $\forall l \in \mathcal{L}$. Clearly, $b \in \bigcap_{l \in \mathcal{L}} \mathcal{B}_l$ but $b \notin \bigcup_{s \in \mathcal{S} \setminus \mathcal{L}} \mathcal{B}_s$. It follows that $b \in \bigcap_{l \in \mathcal{L}} \mathcal{B}_l \setminus \bigcup_{s \in \mathcal{S} \setminus \mathcal{L}} \mathcal{B}_s$. This is true for some

subset $\mathcal{L} \subseteq \mathcal{S}$. Second, we show that this subset is unique. Let again \mathcal{L} be the subset of largest cardinality such that $b \in \mathcal{B}_l$, $\forall l \in \mathcal{L}$ and choose a different subset $\mathcal{V} \subseteq \mathcal{S}$, $\mathcal{V} \neq \mathcal{L}$. Then, either $b \notin \bigcap_{l \in \mathcal{L}} \mathcal{B}_l$ or $b \in \bigcup_{s \in \mathcal{S} \setminus \mathcal{L}} \mathcal{B}_s$. The element b is thus only included in $\bigcap_{l \in \mathcal{L}} \mathcal{B}_l \setminus \bigcup_{s \in \mathcal{S} \setminus \mathcal{L}} \mathcal{B}_s$. ■

By choosing $\mathcal{B}_k = \mathcal{A}_k^n \setminus \bigcup_{k \in \mathcal{S}} \mathcal{A}_k^n$, the result follows by elementary set operations.

ACKNOWLEDGEMENT

The authors would like to thank Prof. Gerhard Kramer for the fruitful and insightful discussions throughout the development of the present work.

REFERENCES

- [1] E. Casini, R. D. Gaudenzi, and O. del Rio Herrero, "Contention Resolution Diversity Slotted ALOHA (CRDSA): An Enhanced Random Access Scheme for Satellite Access Packet Networks," *IEEE Trans. Wireless Commun.*, vol. 6, pp. 1408–1419, Apr. 2007.
- [2] Y. Yu and G. B. Giannakis, "High-Throughput Random Access Using Successive Interference Cancellation in a Tree Algorithm," *IEEE Trans. Inf. Theory*, vol. 53, no. 12, pp. 4628–4639, Dec. 2007.
- [3] S. Gollakota and D. Katabi, "Zigzag Decoding: Combating Hidden Terminals in Wireless Networks," in *Proceedings of the ACM SIGCOMM 2008 Conference on Data communication*, ser. SIGCOMM '08, 2008, pp. 159–170.
- [4] G. Liva, "Graph-Based Analysis and Optimization of Contention Resolution Diversity Slotted ALOHA," *IEEE Trans. Commun.*, vol. 59, no. 2, pp. 477–487, Feb. 2011.
- [5] E. Paolini, G. Liva, and M. Chiani, "High Throughput Random Access via Codes on Graphs: Coded Slotted ALOHA," in *Proc. 2011 IEEE Int. Conf. Commun.*, Kyoto, Japan, Jun. 2011.
- [6] C. Kissling, "Performance Enhancements for Asynchronous Random Access Protocols over Satellite," in *Proc. 2011 IEEE Int. Conf. Commun.*, Kyoto, Japan, Jun. 2011.
- [7] A. Tehrani, A. Dimakis, and M. Neely, "SigSag: Iterative Detection Through Soft Message-Passing," vol. 5, no. 8, pp. 1512–1523, Dec. 2011.
- [8] G. Liva, E. Paolini, M. Lentmaier, and M. Chiani, "Spatially-Coupled Random Access on Graphs," in *Proc. IEEE Int. Symp. on Information Theory*, Cambridge, MA, USA, Jul. 2012.
- [9] C. Stefanovic, P. Popovski, and D. Vukobratovic, "Frameless ALOHA Protocol for Wireless Networks," *IEEE Commun. Lett.*, vol. 16, no. 12, pp. 2087–2090, 2012.
- [10] K. Narayanan and H. D. Pfister, "Iterative Collision Resolution for Slotted ALOHA: An Optimal Uncoordinated Transmission Policy," in *Proc. 2012 7th International Symposium on Turbo Codes and Iterative Information Processing (ISTC)*, Gothenburg, Sweden, Aug. 2012.
- [11] G. L. Choudhury and S. S. Rappaport, "Diversity ALOHA - A Random Access Scheme for Satellite Communications," *IEEE Trans. Commun.*, vol. 31, pp. 450–457, 1983.
- [12] E. Paolini, G. Liva, and M. Chiani, "Graph-Based Random Access for the Collision Channel without Feed-Back: Capacity Bound," in *Proc. 2011 IEEE Global Telecommun. Conf.*, Houston, Texas, Dec. 2011.
- [13] M. Zorzi, "Mobile Radio Slotted ALOHA with Capture, Diversity and Retransmission Control in the Presence of Shadowing," *Wireless Networks*, vol. 4, pp. 379–388, Aug 1998.
- [14] E. Perron, M. Rezaeian, and A. Grant, "The On-Off Fading Channel," in *Proc. IEEE Int. Symp. on Information Theory*, Yokohama, Japan, Jul. 2003, p. 244.
- [15] N. Abramson, "The ALOHA System - Another Alternative for Computer Communications," in *Proc. 1970 Fall Joint Computer Conference*, vol. 37. AFIPS Press, 1970, pp. 281–285.
- [16] A. B. Slomson, *An Introduction to Combinatorics*. Chapman and Hall, 1991.
- [17] D. Slepian and J. Wolf, "Noiseless Coding of Correlated Information Sources," *Information Theory, IEEE Transactions on*, vol. 19, no. 4, pp. 471–480, 1973.
- [18] A. F. Dana, R. Gowaikar, R. Palanki, B. Hassibi, and M. Effros, "Capacity of Wireless Erasure Networks," *Information Theory, IEEE Transactions on*, vol. 52, no. 3, pp. 789–804, 2006.

- [19] T. Ho, M. Médard, R. Koetter, D. R. Karger, M. Effros, J. Shi, and B. Leong, "A Random Linear Network Coding Approach to Multicast," *Information Theory, IEEE Transactions on*, vol. 52, no. 10, pp. 4413–4430, 2006.